

1 In the claims:

2 1. A method for retransmission of transactions in a multi-processor computer
3 architecture, comprising:

4 at a source node in the computer architecture, the source node comprising a
5 retransmit buffer, designating a transaction for transmission to a destination node in the
6 computer architecture, the destination node comprising a receive buffer, wherein the
7 transaction is designated for transmission over a first path in a first flow control class;

8 retrieving a sequence number for the designated transaction;

9 comparing the retrieved sequence number for the designated transaction to
10 sequence numbers in the retransmit buffer, wherein if the comparison does not show a
11 match:

12 attaching the retrieved sequence number to the designated transaction,

13 placing the designated transaction in the retransmit buffer, and

14 sending the designated transaction to the destination node; and

15 wherein if the comparison shows a match, transmitting the designated transaction over a
16 second path.

17 2. The method of claim 1, wherein the designated transaction in the retransmit buffer
18 times out, comprising:

19 retrieving the sequence number of a most recent transaction sent to the destination
20 node along the first path;

21 sending a probe request to the destination node along the second path, the probe
22 request including the sequence number of the timed-out transaction and the sequence
23 number of the most recent transaction;

24 deconfiguring the first path; and

25 updating the sequence number in the retransmit buffer.

26 3. The method of claim 2, wherein the destination node receives the probe request,
27 the method at the destination node, comprising:

28 determining a sequence number for a most recent transaction entry in the receive
29 buffer for a transaction from the source node along the first path and the first flow control
30 class; and

31 determining:

32 (a) if the sequence number for the most recent transaction entry in the
33 receive buffer equals the timed-out transaction sequence number minus one, and

1 (b) if the sequence number for the most recent transaction entry in the
2 receive buffer lies between the timed-out transaction sequence number and the sequence
3 number of the last transaction sent by the source node.

4 4. The method of claim 3, wherein if either (a) or (b) are met:

5 sending a probe response to the source node along the second path;

6 indicating to the source node if the timed-out transaction is received at the
7 destination node; and

8 sending the source node the most recent sequence number acknowledged.

9 5. The method of claim 4, wherein the source node receives the probe response,
10 comprising resuming transmission for all transactions in the retransmit buffer for which
11 an acknowledgement has not been received.

12 6. The method of claim 3, wherein neither (a) nor (b) are met, and wherein an error
13 condition is designated.

14 7. The method of claim 4, further comprising:

15 sending a plunge transaction to the destination node over the first path, the plunge
16 transaction indicating a sequence number for retransmission of transactions should the
17 first path be reconfigured; and

18 updating the sequence number.

19 8. The method of claim 7, wherein the destination node receives the plunge
20 transaction, the method at the destination node, comprising:

21 determining if the first path is deconfigured from the source node; and

22 if the first path is deconfigured from the source node, sending a plunge response
23 to the source node over the first path.

24 9. The method of claim 8, wherein the destination node determines that the first path
25 is not deconfigured, wherein an error condition exists, comprising:

26 informing the source node; and

27 deconfiguring the first path from the destination node.

28 10. The method of claim 8, further comprising:

29 receiving the plunge response; and

30 reconfiguring the first path from the source node.

31 11. The method of claim 2, wherein the pro-

32 determining if a third path is

11 If the third path is available.

deconfiguring the second path, and

updating the sequence number in the retransmit buffer; and

if the third path is not available, designating an error condition.

4 12. The method of claim 1, wherein the destination node receives the designated
5 transaction, the method at the destination node, comprising:

6 determining from the receive buffer a sequence number from a most recent
7 transaction from the source node along the first path in the first flow control class; and

8 comparing the sequence numbers of the designated transaction and the most
9 recent transaction.

10 13. The method of claim 12, wherein the sequence number of the designated
11 transaction is subsequent to the sequence number of the most recent transaction,
12 comprising, determining if the first path is configured.

13 14. The method of claim 13, wherein the first path is configured, comprising:

14 accepting the transaction;

15 updating a receive sequence number in the receive buffer; and

16 sending an acknowledgement to the source node along the first path.

17 15. The method of claim 12, wherein the sequence number of the designated
18 transaction is not subsequent to the sequence number of the most recent transaction,
19 comprising dropping the transaction.

16. An apparatus for retransmission of transactions in a multi-processor computer
architecture, comprising:

22 a source node having a retransmit buffer, wherein the source node stores
23 transactions transmitted from the source node;

24 a send_seqid table comprising a sequence number for each transaction sent from
25 the source node;

26 a destination node comprising a receive buffer, wherein the destination node
27 stores transactions transmitted from the destination node;

28 a receive_seqid table comprising a sequence number for each transaction sent
29 from the destination node, wherein the source node sends normal transactions to the
30 destination node and the destination node sends acknowledgements for the normal
31 transactions to the source node, wherein if the source node does not receive an
32 acknowledgement within a specified time, a corresponding normal transaction in the
33 retransmit buffer times out; and

1 a probe transaction, whereby the source node queries the destination node for a
2 timed-out transaction.

3 17. The apparatus of claim 16, wherein the probe transaction comprises a sequence
4 number of the timed-out transaction and the sequence number of the most recent
5 transaction sent from the source node to the destination node, wherein the source node
6 comprises:

7 means for deconfiguring the first path; and
8 means for updating the sequence number in the retransmit buffer.

9 18. The apparatus of claim 17, wherein the destination node receives the probe
10 transaction, the destination node, further comprising:

11 means for determining a sequence number for a most recent transaction entry in
12 the receive buffer for a transaction from the source node along the first path and the first
13 flow control class; and

14 means for determining:

15 (a) if the sequence number for the most recent transaction entry in the
16 receive buffer equals the timed-out transaction sequence number minus one, and

17 (b) if the sequence number for the most recent transaction entry in the
18 receive buffer lies between the timed-out transaction sequence number and the sequence
19 number of the last transaction sent by the source node.

20 19. The apparatus of claim 18, wherein if either (a) or (b) are met, the destination
21 node:

22 sends a probe response to the source node along the second path;
23 indicates to the source node if the timed-out transaction is received at the
24 destination node; and

25 sends the source node the most recent sequence number acknowledged.

26 20. The apparatus of claim 19, wherein the source node receives the probe response,
27 wherein the source node comprises means for resuming transmission for all transactions
28 in the retransmit buffer for which an acknowledgement has not been received.

29 21. The apparatus of claim 18, wherein neither (a) nor (b) are met, and wherein an
30 error condition is designated.